

F. Chlebana
Nov 1 2006

Status of the CSL Upgrade

I. Bizjak, F.Chlebana, G. Guglielmo, (J.Lee), K.McFarland,
R.Sarkis, JJ.Schmidt, M.Shimajima,
R.Snider, G.Yu, D.Zhang, DataComm, Computing Division

Cast of Characters

Richard Sarkis: *CSL software development*

JJ: *Installation and planning*

Geum bong Yu: *Integration testing*

Ilija Bizak: *Monitoring software*

Tatsuya Masubuchi: *Hardware testing and configuration*

Makoto Shimojima, Kevin McFarland, Jedong Lee: *CSL experts*

Dehong Zhang, Gerald Guglielmo, Rick Snider: *Down stream logging*

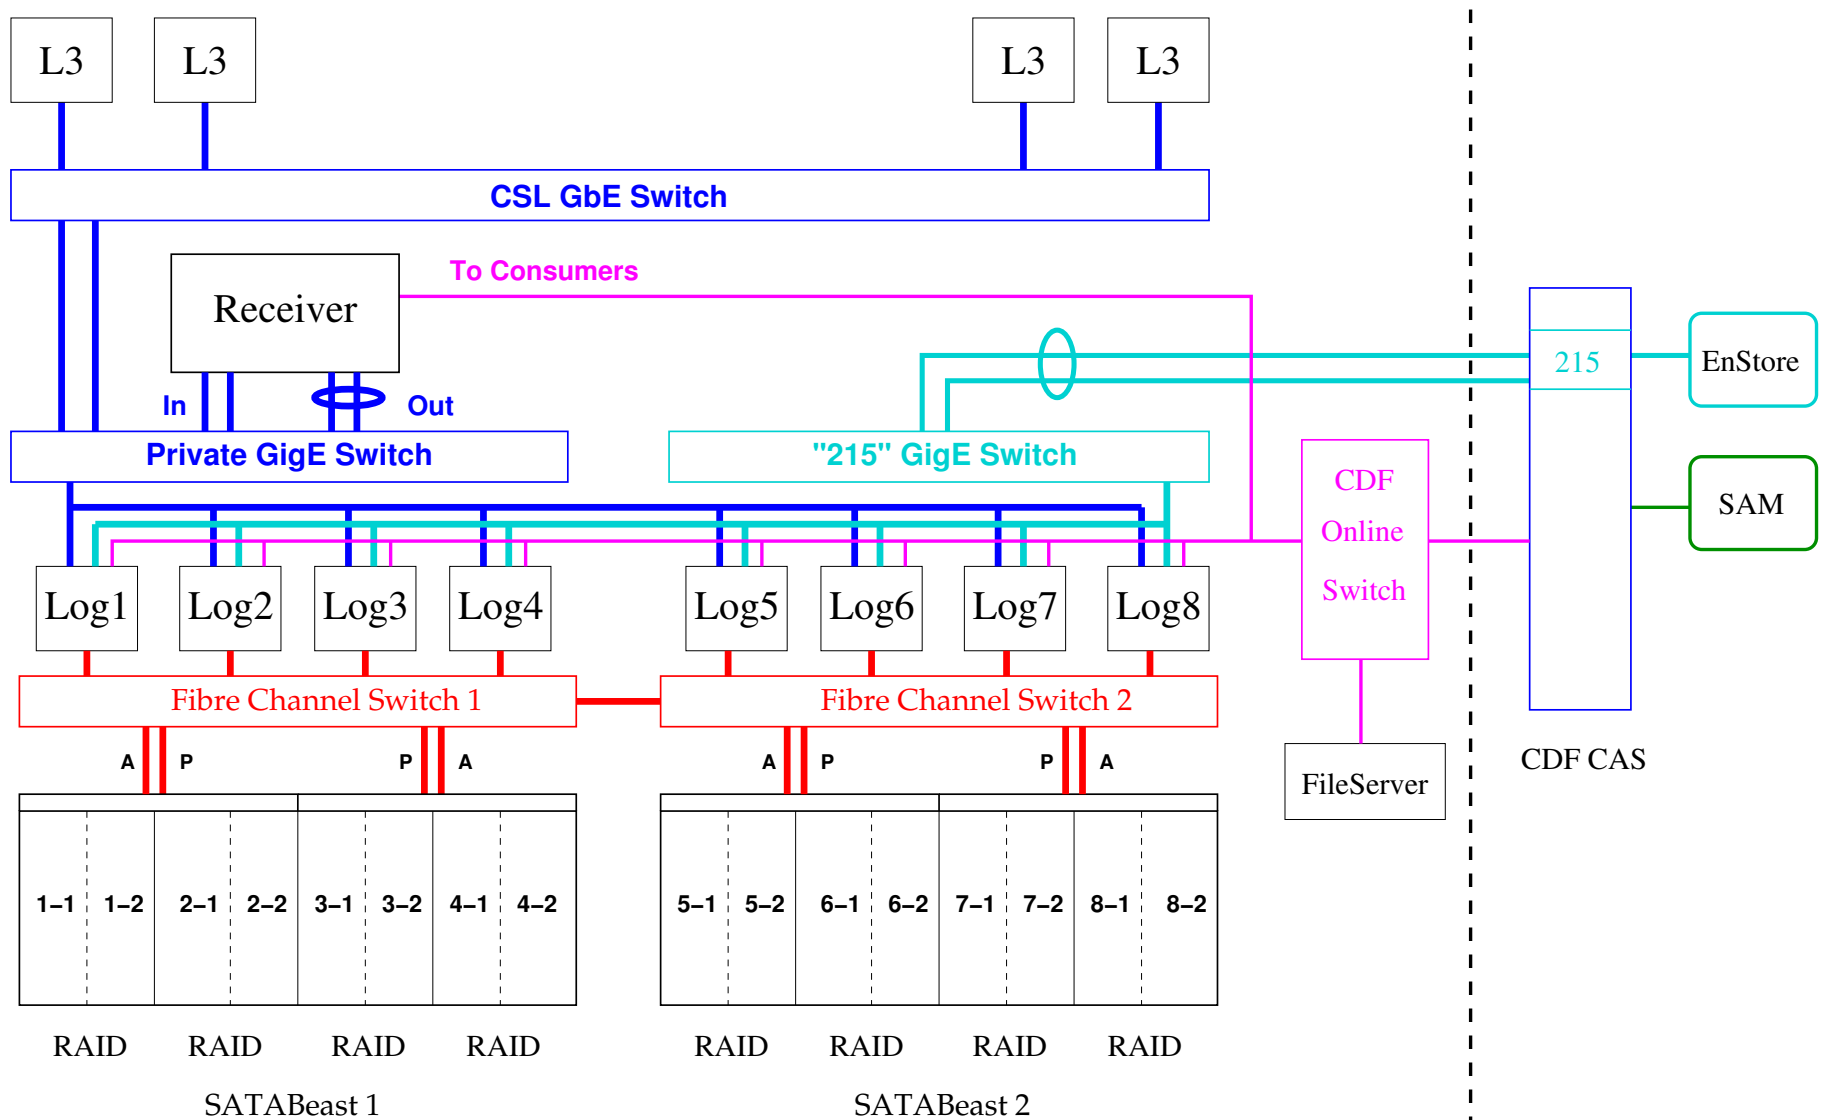
CD/Rex (Jason Hedden): *System configuration*

F. Chlebana: *Scallywag*

Progress Since Last Time

- *Hardware Optimization and Robustness*
- *Data Sorting Test*
- *Orphaned File Checking*
- *Dry Run Test*
- *Commissioning Week*
- *Monitoring Improvements*
- *Documentation and Operating Instructions*
- *Support Issues*
- *Remaining Tasks*

Architecture Overview



Hardware Optimization and Robustness

Were getting reset errors and erratic performance for the disk IO

→ We swapped fibre cables and balanced the load in the NAS

Kafkaesque problems with one of the logger nodes (b0csl26)

Initially had memory related errors

Memory shipped back to D1 but delivered to wrong person

D1 gave us a spare module

Still had problems


Swapped mother board

Still had memory problems

Determined we cannot mix different rank memories

→ *Still need to get correct memory*

Network Tests



Sending data from receiver node to *8 logger nodes*

Bandwidth is equally shared

Using two GigE links *bonded* together




Total network rate: $8 \times 28 = 224 \text{ MB/s}$

Verified that we get the expected full bandwidth when two node send to Receiver on separate input NICs while Receiver sends to two different Logger nodes on bonded output NICs

Disk IO Tests

Concurrent Read/Write tests on *all* 8 logger nodes

Bandwidth has been balanced now see the same IO performance across the 8 logger nodes.




Write = 57 + 56 + 55 + 55 + 56 + 56 + 55 + 55 = 445 MB/s

Read = 75 + 75 + 72 + 73 + 75 + 75 + 73 + 72 = 590 MB/s

Ran extended concurrent disk IO tests on all 8 logger nodes at the same time.

Tests robustness of the NAS and controllers on the SATABeast



No problems seen!

This really loads the system far beyond we can hope to achieve in normal operations.

Data Sorting Test

The CSL does not touch the data, it directs the data to the designated output stream and to the consumers based on the *event header* that is passed to the CSL from L3.

Repeated the same data sorting tests as done with the old CSL.

The data sorting test uses a software sender process (SA_sender) that plays a known output pattern through the CSL.

We then verify that the expected number of events end up in the designated stream.

Events just flat out going to the wrong stream

Tested with SA_sender which has csl streammask argument.

By increasing the streammask one by one up to 17, I checked the output csl files. Here's the streammask-output stream file matching.

```
1  - A,  
2  - B, 3  - A+B  
4  - C, 5  - A+C, 6  - B+C, 7  - A+B+C  
8  - D, 9  - A+D, 10 - B+D, 11 - C+D, 12 - A+B+D, 13 - A+C+D, 14 - B+C+D,  
    15 - A+B+C+D  
16 - E, .....
```

Verified that the number of events recieved in the output streams were as expected.

Events with “no marked stream” being mishandled (logged) or not being made available to consumers

Tested with SA_sender.

Verified that events with “no marked stream” do not end up in data files.

Events selected for multiple streams being mishandled. For example: going to no stream, going to only one of the selected streams, etc.

Tested with SA_sender.

The received number of events were same as expected.

The event structure is tested by offline processing and by the consumers.

→ *We verified that a cosmic run taken during the Commissioning week could be handled offline. The run was successfully processed.*

→ *We verified that the consumers were able to accept events.*

→ *Ran through series of calibrations*

Orphaned File Checking

We want to make sure that we do not have files left over on the disk buffer that are not accounted for.

Files are being kept track of on the “CSL Fetch” web page

Node	Updated	FSS	Lun	CSL			Stager		
				Open	To Store	Fetch Err	New	Queued	Submitted
b0csl21	<u>1030</u> <u>14:05</u>	<u>Alive</u>	1	<u>0</u>	0	<u>0</u>	<u>0</u>	<u>0</u>	<u>0</u>
			2	<u>0</u>	0	<u>0</u>	<u>0</u>	<u>0</u>	<u>0</u>
b0csl22	<u>1030</u> <u>14:05</u>	<u>Alive</u>	1	<u>15</u>	0	<u>0</u>	<u>0</u>	<u>0</u>	<u>0</u>
			2	<u>0</u>	0	<u>0</u>	<u>0</u>	<u>0</u>	<u>0</u>
b0csl23	<u>1030</u> <u>14:05</u>	<u>Alive</u>	1	<u>5</u>	0	<u>0</u>	<u>0</u>	<u>0</u>	<u>0</u>
			2	<u>3</u>	0	<u>0</u>	<u>0</u>	<u>0</u>	<u>0</u>

Requires expert interpretation

Working on automating the check

Monitor files when they are first written to the buffer disk (.open) and at the point that they are written to tape.

→ *List of files generated by the logger process at the first place data is written to disk.*

→ *Later will be checked against a list of files written to tape by the Stager process*

First version implemented... but needs to be able to handle more realistic operating conditions.

→ *Renaming file based on lowest run section*

→ *Not all files will be sent to tape*

Orphaned File Checking

We do see a couple of files that are left in the “.open” area.

This is the directory that the CSL Logger first writes to. When the file is closed it is moved up one directory and handed over to the Stager process.

The current monitoring is able to alert us about this issue.


Other problems at the time and it looks like the run was abnormally aborted...

Being investigated...

Dry Run Tests

Prior to using the system in operations we ran an extended test.

L3Sender → CSL Receiver → Distributed to 8 logger nodes → written to disk → Eight Stager processes transfer data to mover nodes in Feynman → Data written to tape!




We had a “dry run” test using the software event sender. Data was sent through the full CSL chain at about 100 MB/s continuously for more than 24 hours.

Revealed some interesting features which were addressed...

- *IO was not balanced (introduced rate limit in system...)*
- *Verified throttling problem was gone*
- *Enstore mover node(s) “got stuck”...*
- *Worked through details on exception handling...*

The tape logging rate is able to keep up with the data buffering rate (that is why you see the periodic pattern...)



More details at:

<http://ncdf76.fnal.gov/~chlebana/daq/cslUpgrade/commissioning/dryRun/>

In our testing we really stressed the system to the limit and improved the performance and reliability.

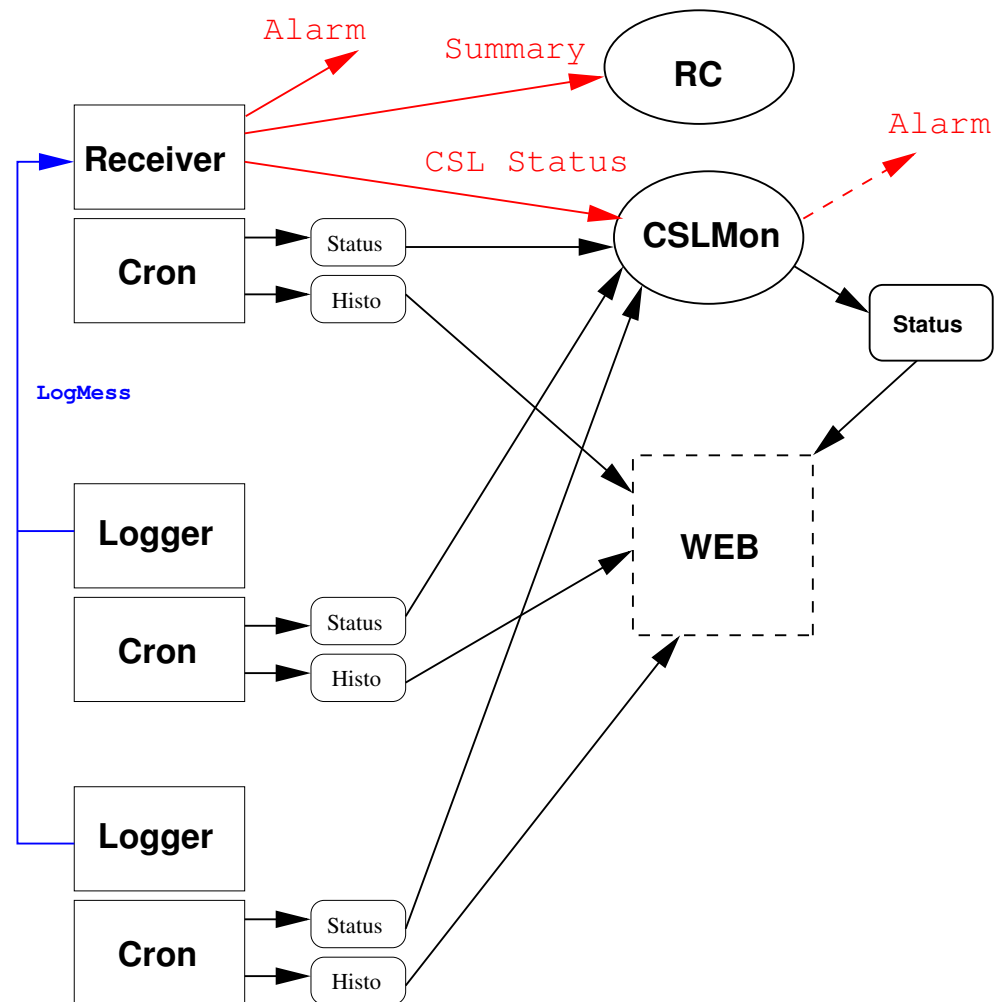
During high rate testing revealed problems in the software that is already being used at D0.

Adapting solutions and unifying software

Monitoring

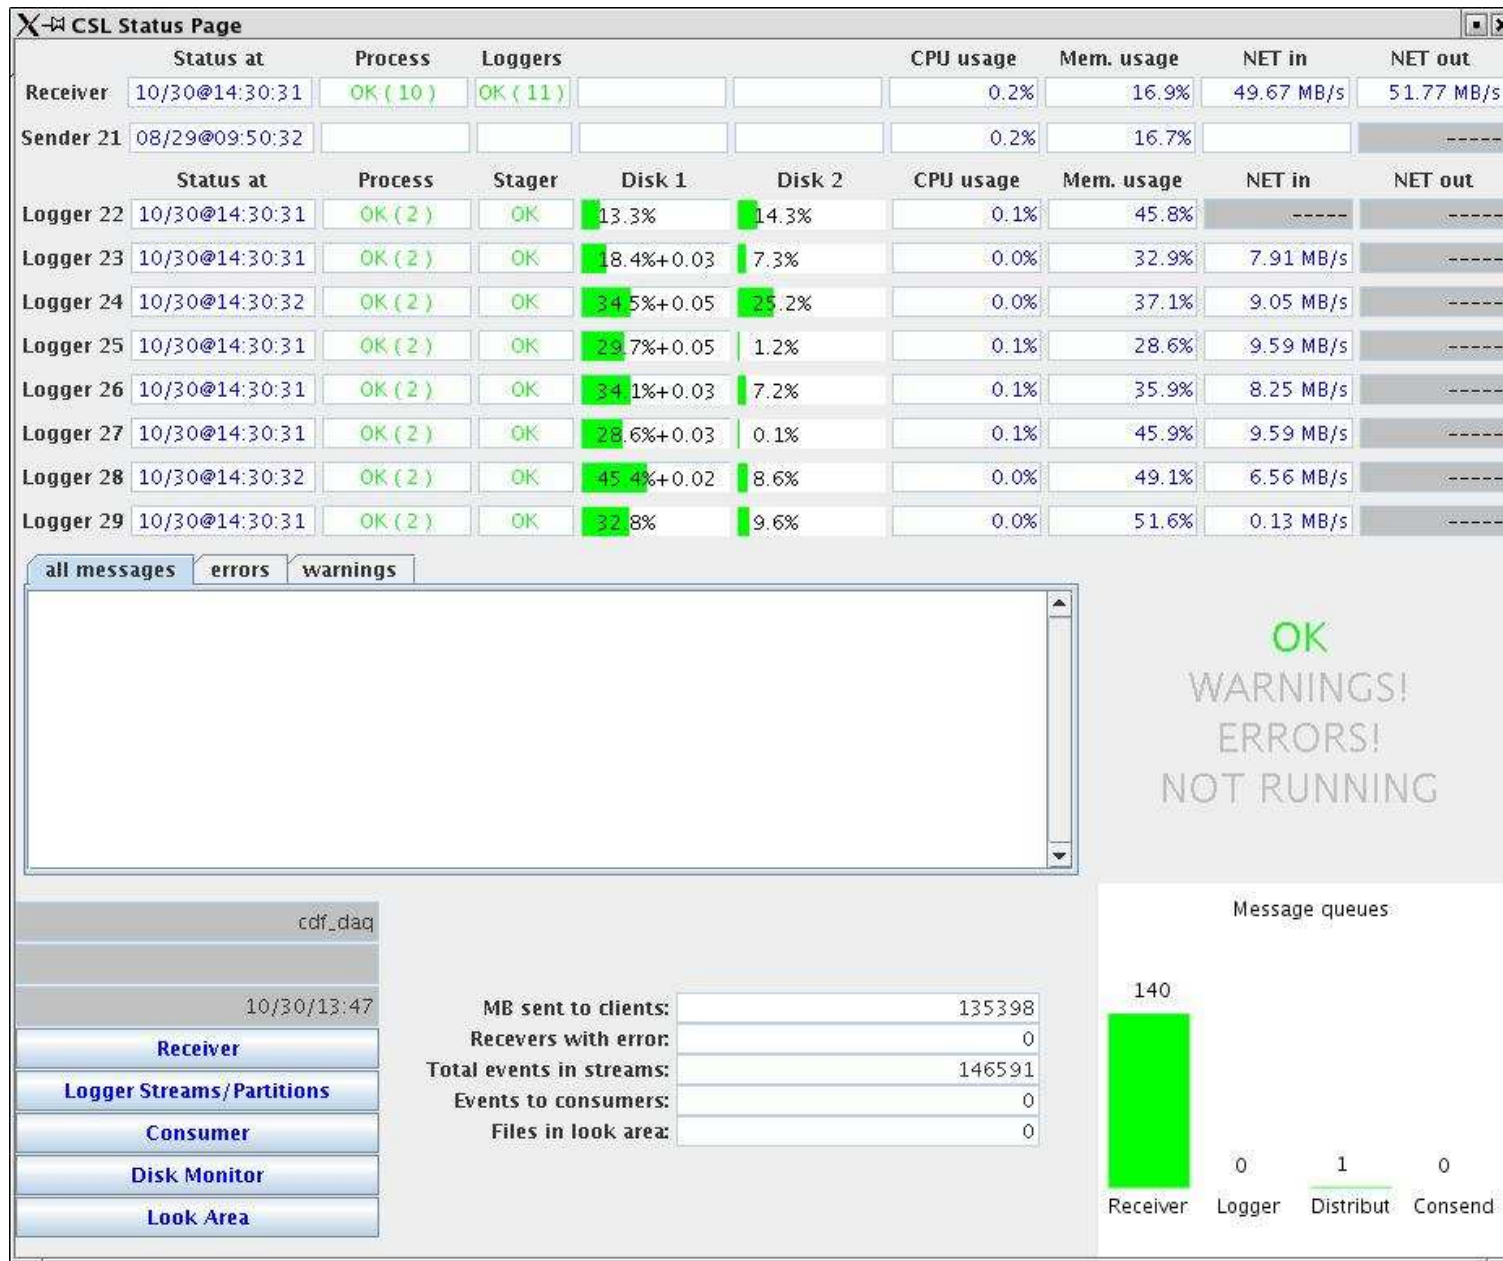
Content of messages will mostly remain the same, sources of information will change...

- *Logger processes send "LogMess" to Receiver*
- *Receiver broadcasts "CSL Status" message as before*
- *CSLMon listens for message, interprets and display statistics*
- *Logger can provide additional information as text files*
- *Content formatted for display on the web*



Messaging updated to use smart sockets!

The New CSLMon...



Most important information is grouped on the main page.

Shift crew does not have to open additional windows
(Avoids information overload...)

Same monitoring information is now available in the new CSL.

NET in	NET out
49.67 MB/s	51.77 MB/s

NET in	NET out
-----	-----
7.91 MB/s	-----
9.05 MB/s	-----
9.59 MB/s	-----
8.25 MB/s	-----
9.59 MB/s	-----
6.56 MB/s	-----
0.13 MB/s	-----

Detail showing the total logging rate (51 MB/s) and individual logger node rates.

Other monitoring windows have the same look and feel as before.

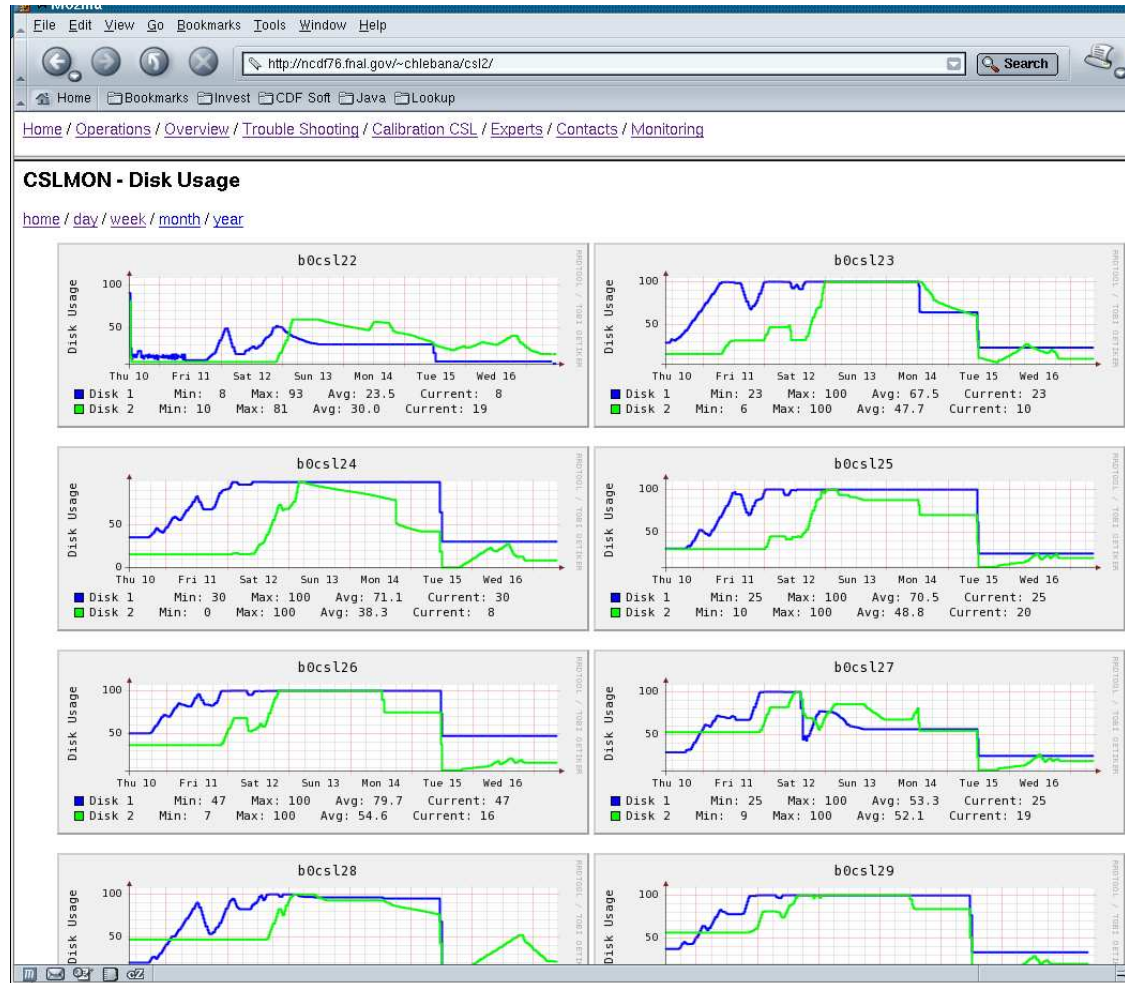
- Receiver
- Logger Streams/Partitions
- Consumer

rates for logger streams and partitions																	
File Info		Logger Rates for partitions [0-16] and streams [0-10]															
		Rates shown: Most recent Rates (MB/s)										Stream definitions					
	Node	Part 0	Part 1	Part 2	Part 3	Part 4	Part 5	Part 6	Part 7	Part 8	Part 9	Part 10	Part 11	Part 12	Part 13	Part 14	Part 15
Run Number		226453	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
stream A	b0csl22-private	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
stream B	b0csl23-private	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
stream C	b0csl24-private	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
stream D	b0csl25-private	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
stream E	b0csl26-private	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
stream F		0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
stream G	b0csl27-private	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
stream H	b0csl28-private	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
stream I	b0csl26-private	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
stream J	b0csl29-private	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Part totals		83.191	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
		Total				83.191				Overall Logging counter				-----			
		Show peak rate		Show average rate		Show current rate		Show number of events		Show event rate							

Logger Streams/Partitions window showing rate information per stream

History monitoring plots available on the web

- Using rrdtools to generate plots
- Using ganglia for system monitoring



<http://www-cdfonline.fnal.gov/ops/ace2help/csl2/>
→ *Monitoring tab*

Monitoring data flow at each step...

The NEW CSL ACE Help Pages - SeaMonkey

Overview Operations Calibration CSL Trouble Shooting Expert Pages Contacts Monitoring

CSL To Enstore Status - Fri Oct 27 10:30:35 2006

SAM DB Server (rawdata_prd) [Alive](#)
 Quota Home / Data / Spool - sam (%) [33 / 12 / 0](#)
 Quota Home / Data / Spool - stager (%) [22 / 57 / 0](#)

Node	Updated	FSS	Lun	CSL			Stager			In FSS	Recent Stores No LastDone	Clsum Warn/Err/Alarm HeartBeat	Clmeta Warn/Err HeartBeat	Lookarea		
				Open	To Store	Fetch Err	New	Queued	Submitted					New	Copied	Error
b0csd21	1027 10:30	Alive	1	24	0	0	0	0	0	0	0 / 0 / 0 1027 10:29	0 / 0 1027 10:30	0	0	0	
			2	0	0	0	0	0	0	0	0	0	0	0		
b0csd22	1027 10:25	Alive	1	15	0	0	0	0	0	0 1026 21:49	0 / 0 / 1 1027 10:27	0 / 0 1027 10:24	0	0	0	
			2	0	0	0	0	0	0	0 0824 10:59	0	0	0			
b0csd23	1027 10:30	Alive	1	5	0	0	0	0	0	0 1020 21:24	0 / 0 / 0 1027 10:29	0 / 0 1027 10:29	0	0	0	
			2	2	0	0	0	0	0	0 0906 20:39	0	0	0			
b0csd24	1027 10:30	Alive	1	26	0	0	0	0	0	0 1020 21:24	0 / 0 / 0 1027 10:29	0 / 0 1027 10:29	0	0	0	
			2	0	0	0	0	0	0	0 0912 03:26	0	0	0			
b0csd25	1027 10:30	Alive	1	35	0	0	0	0	0	0 1020 21:25	0 / 0 / 0 1027 10:27	0 / 0 1027 10:29	0	0	0	
			2	4	0	0	0	0	0	0 0823 18:01	0	0	0			
b0csd26	1027 10:30	Alive	1	48	0	2	0	0	0	0 1026 20:58	0 / 0 / 0 1027 10:28	0 / 0 1027 10:29	0	0	0	
			2	0	0	0	0	0	0	0 0823 16:35	0	0	0			
b0csd27	1027 10:25	Alive	1	29	0	0	0	0	0	0 1020 21:25	0 / 0 / 0 1027 10:27	0 / 0 1027 10:24	0	0	0	
			2	0	0	0	0	0	0	0 0817 17:40	0	0	0			
b0csd28	1027 10:30	Alive	1	38	0	8	0	0	0	0 1020 21:24	0 / 0 / 0 1027 10:27	0 / 0 1027 10:29	0	0	0	
			2	0	0	0	0	0	0	0 0823 18:01	0	0	0			
b0csd29	1027 10:30	Alive	1	22	0	0	0	0	0	0 1020 21:24	0 / 0 / 0 1027 10:29	0 / 0 1027 10:29	0	0	0	
			2	0	0	0	0	0	0	0 0912 10:09	0	0	0			

Consumers

Tested sending data to all consumers during Commissioning Week

In the current system data to the consumers is sent via one GigE link. *Typical rates of 6 MB/s (5% of available bandwidth).*

In the new system we will also use a one GigE link

→ *Same effective bandwidth*

Consumer data path is over the 236b subnet

→ *Cannot saturate the link*


Ultimately limited by processing time in the consumer...

→ *May need to optimize consumer in order to reduce processing time if interested in increasing the event rate.*

Connected consumers and rate information

[illegible]

Commissioning Week



*How can we possibly test
the New CSL when there
is Beam in the Machine...*

an opportunity presented itself...

Commissioning Week

Started using the New CSL during normal operations on Oct 14.

Ran for a full week.

Found some operational issues...

1) Problems when changing Run Configurations

→ Handled by developing a quick recovery procedure

→ Eventually the problem was “fixed”.

2) Look area was filling up with every file

→ The problem was fixed.

Used the “CSL Torture” trigger table for extensive testing...

At the end of the week we worked through most issues and are now comfortable with using the New CSL for data taking.

Commissioning Week

Cosmic data was processed offline

Hi Frank,

We had no problems or remaining issues with accessing and processing raw data taken with the new CSL (run=225118) and sign off on this.

Thanks,

Igor and Guram

G. Pauletta was having problems reading Laser Calibration files

Howard Budd was able to read the files:

These are the laser runs. I can read all the runs that exist since Oct 14. If there is no comment after the run it means its OK, ie I didn't find a problem. (Note 1 channel appears low in 2 runs, but that is not a problem with the file.

→ *Should try recompiling the code with a new software release, he was using 5.3.1.*

→ *Need to follow up on this...*

Additional Development

There are a couple of **non-critical** outstanding issues:

- Want to get “ether channeling” to work between L3 Output and CSL Data switch, *boosts bandwidth from 120 to 240 MB/s.*
- Install memory on b0csl26, *have 2 GB want 4 GB*
- Test UPS software, *gracefully shut down servers*
- Swap in new Fibre cables (happened today), *using borrowed cables*
- Configure spare logger node *will directly connect it to the dotHill array*
- Eventually write data to the new tape robot *In 2007...*

Spare Logger Node Configuration

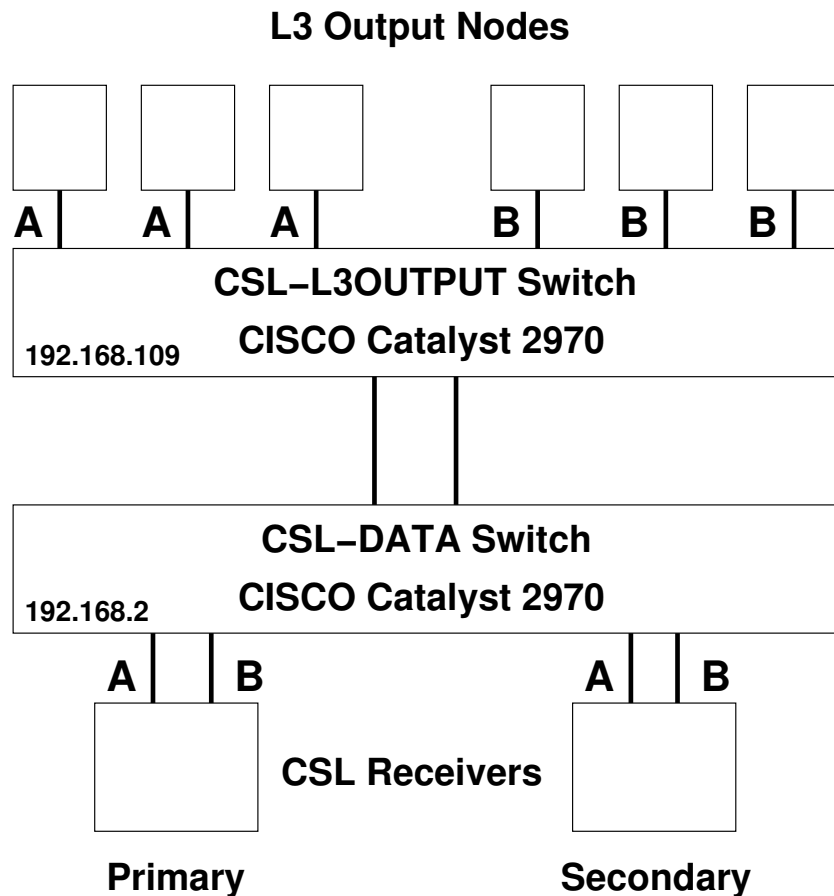
Initially we were going to use b0csl21 as a spare logger node and reconfigure it to connect to the SATABeast

In order to simplify the switch to the spare logger node in case of problems we will connect b0csl21 directly to the dotHill disk array

→ *b0csl21 will always be available - no need for system level reconfiguration, simplifies disaster recovery...*

→ *Can use b0csl30 and b0csl21 as our development platform*

EVB Output/CSL Data Switch



We are planning on using ether channelling on two links between the CSL-L3Output switch and the CSL-Data switch

This will distribute the load across the L3 output nodes and gives us an effective bandwidth of 240 MB/s

Currently only one link is being utilized...

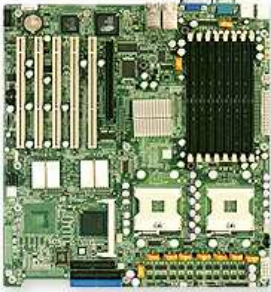
Documentation on the *Internets*

The NEW CSL ACE Help Pages - SeaMonkey

Overview Operations Calibration CSL Trouble Shooting Expert Pages Contacts Monitoring

Overview / Hardware / Performance

Servers



The CSL servers use the [X6DHE-XG2](#) motherboard.


- Dual 3.4 GHz Xeon
- 4 GB memory
- 5 PCI-X slots

There are 8 designated Logger nodes (b0csl20 - b0csl29) and one Receiver node (b0csl20). We have two hot spares, one configured as a spare Logger (b0csl21) and a second configured as a spare Receiver (b0csl30). The hardware of all the servers is the same.

The Reciever has 4 NICs and 1 RAID controller

The Loggers have 2 NICs and 1 HBA

SATABeast




The storage arrays are the SATABeast from [NEXSAN](#).

- Dual controller with 2 ports per controller
- Using 28 500 GB disk.
- Total disk slots available is 42

The SATABeasts are partitioned into 4 seperate RAID5 areas with 2 LUNs/RAID. Traffic to the RAID areas is distributed between the two controllers. Normally one Logger node will mount 2 LUNs belonging to the same RAID.

There are two spare disks which can be used to rebuild any of the 4 RAID.

[QLogic SANbox 5200 Fibre Channel Switch](#)



Done

Support Issues

System Redundancy

Avoid having a single point of failure

- *Multiple fibre paths*
- *RAIDed boot disk*
- *Fully configured hot spares*
- *Available spare parts*

Self repairing, receiver can redirect traffic to free logger/filesystem

System can be reconfigured to bypass failed component

Deep buffers (3 days at 80 MB/s) allow plenty of time to react to down stream problems

Expect most problems can be handled by the CSL pager carrier

Long Term Support Issues

Rochester

- 1 Onsite Graduate Student
- 1 Remote Computer Professional (part time)

Tsukuba

- 1 Onsite Graduate Student

University College London

- 1 Remote PostDoc (part time)

Computing Division

System hardware support CD/Rex

Stager/Logger

Data handling group

Not sufficient to provide ongoing support

What we need:

SPL

Main contact to address operational issues

Monitor performance/elds for potential problems

Coordinate activities

CSL pager carriers

Should have at least 3

System support

CD/Rex - *I think this should be ok...*

CSL Software support

No resident expert

Will be difficult to get quick turn around to address problems

→ *Needs to work when delivered*

→ *Must be available in case of urgent problems*

Monitoring software support

No resident expert

Will be difficult to get quick turn around to address problems

→ *I expect that this will be the place we will continued development as we gain additional operational experience*

Stager/Logger pager coverage

Will need point of contact

Next Steps...

Basic logging software is working and is robust enough to be used in production - *but can always make improvements...*

→ *SPL should be directly involved in updating documentation and monitoring performance during the initial stages of operations*

→ *Good to get a fresh view and excellent opportunity to learn the details of the system*

We have the monitoring framework in place - but this is an area that will require ongoing improvements

→ *Need to develop resident expertise for the monitoring software*

Start using it!!!

We switched to the New CSL on Oct 31... if all goes well, plan to leave it in when we get beam again.

Summary

Hardware Optimization and Robustness

- Working well
- Performing well above the target of 80 MB/s

Data Sorting Test

- Repeated the same data integrity test using for the old CSL

Orphaned File Checking

- Added file monitoring from the first time data is written to disk to when it is written to tape
- Automating the check

Dry Run Test

- Performed extended high rate testing

Commissioning Week

- Used the New CSL for a full week and worked through operational problems
- Ready for use...

Documentation and Operations

- Basic documentation in place
- *Needs updating*

Support Issues

- *Need additional help for ongoing support*

Remaining Tasks

- *There are a couple of non-critical tasks that need to be finished up.*

